

SYNCHRONOUS PEER-TO-PEER MULTIPOINT DATABASE
SYNCHRONIZATION

BACKGROUND OF THE INVENTION

5 [0001] The present invention generally relates to database synchronization and more particularly to a system and method of database synchronization over a peer-to-peer network.

[0002] With the popularity of handheld computing devices (i.e., PDA's, cell phones, etc...) increasing, there is becoming a greater need and ability to share information between devices. Computing networks can be established between the devices so that collaborative information can be shared. The computing devices can form peer-to-peer networks between one another such that information can be shared without
10 the use of a central server to store a database of information.

[0003] Currently, in the prior art, computers use databases stored on a central database server in order to communicate information. Database systems have replication and synchronization capabilities in order to update information on client
15 systems. These synchronization capabilities are usually restricted to simple two-point exchanges between clients and servers. The synchronization and replication capability require the database server to effect dataset reconciliation between multiple users.

[0004] When multiple users without access to the database server wish to
20 synchronize their databases, it must be done by a series of 2-way exchanges. For instance, information must flow in both directions between users in order to synchronize the dataset. At the very least, this involves $2*n$ combinations of

[0005] U.S. Patent No. 6,295,541, entitled "Systems and Methods for Synchronizing Two or More Datasets" describes a method whereby a reference database is maintained which the other devices synchronize to when available. The patent describes a system whereby devices synchronize to the reference database at
10 different times when they are online such that synchronization occurs serially between devices. Furthermore, only one device is considered the reference database such that only one device controls the synchronization process.

BRIEF SUMMARY OF THE INVENTION

2

changes, and send those changes to the initiator. The initiator then decompresses their changes and replicates them into its local database. The initiator then extracts all changes (including those received from other group members) from the local database. Next, the changes are compressed and sent to the other users in the peer-to-peer network. Finally, the changes are decompressed and replicated on the database of each of the other users. In this respect, the databases of each group member will contain all changes from the databases of all the group members.

[0008] The above description describes a bilateral synchronization, in that changes are both sent and received from each peer. A unilateral synchronization is also possible, which consists of either the first half (receiving changes) or the second half (sending changes) of the above description. It will be recognized that the process can be initiated and performed by any user in the peer-to-peer network. In the preferred embodiment of the present invention, the transfer of the changes between the multiple users is done in parallel, although it may possibly be sequential. The transfer can be performed over a wired or wireless network.

[0009] In accordance with the present invention, there is also provided a system for synchronizing databases of multiple users. The system includes a plurality of computers in a peer-to-peer network wherein each computer has a database and software configured to provide synchronization. The software is configured to extract changes from the local database of each of the initiator's peers. The peers' changes are sent to the initiator in parallel as they finish being extracted. Each peer's changes are replicated into the initiator's database. Then all changes (including those received from the peers) are extracted from the initiator's database. Next, the software sends the changes in parallel to the other users of the peer-to-peer network. Finally, the changes are replicated by the software on each database of the other users in order to synchronize the databases.

BRIEF DESCRIPTION OF THE DRAWINGS

[0010] These as well as other features of the present invention will become more apparent upon reference to the drawings wherein:

5 [0011] Figure 1 is a flowchart illustrating a method of unilateral database synchronization from a target computer to an initiating computer according to the present invention;

[0012] Figure 2 is a diagram illustrating unilateral database replication from multiple target computers to the initiating computer;

10 [0013] Figure 3 is a diagram illustrating unilateral database replication from the initiating computer to multiple target computers;

[0014] Figure 4 illustrates the synchronization process on the initiating computer;

[0015] Figure 5 is a diagram illustrating bilateral database replication between target computers and the initiating computer;

15 [0016] Figure 6 is a diagram illustrating the selection of database replication techniques by a user; and

[0017] Figure 7 is a flowchart illustrating a method of unilateral database synchronization from the initiating computer to the target computer.

DETAILED DESCRIPTION OF THE INVENTION

[0018] Referring now to the drawings wherein the showings are for purposes of illustrating a preferred embodiment of the present invention only, and not for purposes of limiting the same, Figure 2 shows unilateral database replication
5 between an initiating computer 10 and target computers 12a, 12b and 12c according to the present invention. Both the initiating computer 10 and the target computers 12a, 12b, and 12c are computing devices having a memory and capable of storing and processing information. For instance the computing devices may be PDA's (i.e., Personal Digital Assistants) which are in electronic communication with one another
10 through a peer-to-peer network. The network may be a wireless network such as Bluetooth, or could be a wired network such as an Ethernet. In either case, the peer-to-peer network provides electronic communication between the initiating computer 10 and the target computers 12 without the use of a server.

[0019] The initiating (source) computer 10 is synchronized with the other computers
15 12 by the transfer of database information from the target computers 12 to the initiating computer 10. As seen in Figure 2, each of the target computers 12 transfer information to the initiating computer 10 in the same manner. Accordingly, the description below will be for the transfer of database information from target computer 12a to source computer 10, yet the process is identical for the transfers
20 between target computers 12b and 12c to source computer 10. As will be further explained below, the processing and transfers may occur in parallel for each of the target computers 12a, 12b, and 12c to the initiating computer 10.

[0020] The process for synchronizing the target computers 12 to the initiating computer 10 begins by the initiating computer 10 sending out a synchronization
25 request to the target computers 12. The synchronization command informs the target computers that the initiating computer 10 wishes to synchronize databases with them. It will be recognized that any computer in the network may be the

initiating computer such that it is possible for any computer to start the synchronization process. The term initiating computer designates the computer that wishes the synchronization process to begin.

5 [0021] Referring to Figures 1 and 2, in step 100, the target computer 12a has a source database 14a that is created in memory (i.e., harddrive, RAM, etc...). The source database 14a contains information which is shared between the target computer 12a and the other computing devices (i.e. computers 12b, 12c and 10). Next, in step 102, any changes to the source database 14a are extracted to create a temporary extracted database 16a containing at least all relevant changes to both the
10 target computer 12a and the initiating computer 10. The extracted database 16a corresponds to the replica source 14a stored on the target computer 12a. All changes to the source database 14a from a user-specified date and time, or from the last time a synchronization occurred, are extracted and copied to the extracted database 16a. Accordingly, as seen in step 104, the extracted database 16a is created
15 by extracting and copying the changes to the source database 14a. The extraction and copying is performed using the database API of the target computer 12a if such an API is available. In the case where no such API is available, changes are tracked as data is written to each database and extracted by the synchronization system itself.

20 [0022] Next, the extracted database 16a is compressed in step 106 to create a compressed database 18a in step 108. The compressed database 18a is created by using well known compression techniques on the extracted database 16a. Once the compressed database 18 is created, then the extracted database 16a is expunged in step 109.

25 [0023] Referring to step 110, the compressed database 18a is transferred to the initiating computer 10. Specifically, the file of the compressed database 18a is transferred to the initiating computer 10 in response to a request using any well

known file transfer technique over any type of network, as previously described. Each of the target computers 12 will transfer information to the initiating computer 10 in parallel or in sequence when it is not possible to perform parallel communications. Typically, the compressed database 18a is transferred to a temporary transferred database 20a that is created on the initiating computer 10 in step 112. Once the transfer is complete, in step 113, the compressed database 18a on the target computer 12a is expunged.

[0024] Once the transferred database 20a has been created on the initiating computer 10, the transferred database 20a is synchronized with the source database 24 of the initiating computer 10. Specifically, in step 114 of Figure 1, the transferred database 20a is decompressed on the initiating computer 10, as seen in Figure 4. The transferred database 20a is decompressed using a complementary decompression technique to that of step 106. Accordingly, in step 116, a decompressed database 22a is created on the initiating computer 10. Once the decompressed database 22a is created, then the transferred database 20a is expunged in step 115.

[0025] The decompressed database 22a is then replicated with the source database 24 of the initiating computer 10 in step 118. This operation is performed using the database API of the initiating computer 10 if such an API is available. In the case where no such API is available, replication is performed by the synchronization system itself. The process of replication causes the changes copied in the extracted database 16a to be incorporated into the source database 24 of the initiating computer 10. Once the decompressed database 22a is replicated onto the source database 24, the updated source database 24 is created in step 120 which has the changes and is identical to the source database 14a. Finally, the decompressed database 22a is expunged in step 122.

[0026] Each time a target computer 12 sends changes to the initiating computer 10, the synchronization process previously described (i.e., transfer, compression, decompression and replication) are completed in full before another synchronization from another target computer 12 is processed.

5 [0027] The above-described procedure is operative to send changes from the target computers 12 to an initiating computer 10. This procedure typically occurs when an initializing user wishes to receive changes from the other computers. However, the procedure can also be used if the initializing user wishes to transfer changes to other computers. In that instance, changes from the database of the initiating computer 10
10 would be transferred to the other computers 12a, 12b, and 12c.

[0028] Referring to Figure 3, a diagram for the unilateral synchronization of multiple computers from the initiating computer 10 to target computers 12 is shown. Synchronization between the initiating computer 10 and target computers 12 occurs in parallel. Each target computer 12 has a transferred database, a decompressed
15 database, and a target database that are created during the process of synchronization.

[0029] Referring to Figures 3 and 7, the process for synchronizing the target computers 12 to the initiating computer 10 is similar to the process of synchronization described in Figure 1. For example, the process begins with the
20 initiating computer 10 sending a synchronization request to the target computers 12a, 12b, and 12c informing them that the initiating computer 10 wishes to send them changes to their databases. The synchronization process proceeds according to Figure 7 such that changes in the source database 24 of the initiating computer 10 are extracted in step 702 to create an extracted database 26 in step 704. In step 706,
25 the extracted database 26 is compressed to create a compressed database 28 in step 708. Once the compressed database 28 is created, the extracted database 26 is

expunged in step 709. The compressed database 28 is transferred to the target computers 12a, 12b, and 12c in parallel over the computer network.

[0030] Each of the target computers 12a, 12b, and 12c receives the compressed database 28 and creates a respective transferred database 30a, 30b, and 30c. For simplicity, the following description will be for target computer 12a. However, it will be recognized that the following synchronization process occurs in any of the target computers (i.e., 12b, and 12c) synchronizing to the initiating computer 10. After the compressed database 28 is transferred and the transferred database 30a is created on the target computer 12a, then the compressed database 28 on the initiating computer 10 is expunged in step 713. In step 714, the compressed database 30a on the target computer 12a is decompressed to create a decompressed database 32a in step 716. The transferred database 30a is then expunged in step 715. The changes from the source database 24 are then replicated onto the database 14a of target computer 12a in step 718. In this regard, the database 14a will be updated with the changes from initiating computer 10, as seen in step 720. Finally, the decompressed database 32a is expunged from the target computer 12a.

[0031] The above-described method is concurrently performed on each of the target computers 12a, 12b, and 12c such that synchronization occurs simultaneously. Accordingly, the synchronization process of the present invention is a parallel method whereby each of the target computers 12 can synchronize with the initiating computer 10 quickly.

[0032] A bilateral database replication between an initiator and n peers consists of first n unilateral replications conducted in parallel to replicate all of the peers' changes into the initiator's database, followed by another n unilateral replications, again conducted in parallel, to replicate the accumulated changes from the initiator's database into the peers' local databases. Referring to Figure 5, the initiating computer 10 receives changes from each of the target computers 12 in a sequential

manner through the process described for Figure 2. While in no defined order, each transfer and synchronization is completed before the next transfer and synchronization is started. For example, the synchronization process from target computer 12a to the transferred database 20a and synchronization to source database 24 will be completed before the next synchronization process from either target computer 12b or 12c. Once all of the target computers 12 have synchronized their changes to the source database 24, then the changes will be sent back to all of the target computers 12 in order to fully synchronize each target computer 12 to one another. Specifically, the source database 24 is now updated with the changes from each of the target computers 12. The initiating computer 10 will send all of its changes to the target computers 12 through a parallel manner as described for Figure 3. The updated source database 24 is sent to each target computer 12 and each target computer 12 synchronizes the changes into its own database 14 in parallel. Accordingly, complete synchronization of all databases (i.e., initiating computer 10 and target computers 12) occurs in a two-stage process (bi-lateral exchange).

[0033] Even though the above-mentioned synchronization process for multiple computers has been described as being between an initiating and target computers, it will be recognized that any computer in the peer-to-peer network can begin the synchronization process. The computer that initiates the process will push and pull the transfer of files as needed. Furthermore, during multi-user synchronization, if any computer (other than the initiator) disappears from the network or cancels the process, the synchronization process can still proceed between the other computers.

[0034] Referring to Figure 6, a flowchart showing possible synchronization processes is shown. As previously mentioned, either bilateral or unilateral synchronization between the initiating computer 10 and the target computers 12 can occur. The user determines the type of synchronization process in step 600 of Figure 6. The user can send changes from the initiating computer as described for

Figure 3, receive and send changes as described for Figure 5, or receive changes as described for Figure 2. If the user decides to send changes only from the initiating computer 10 to target computers 12 (i.e., unilateral exchange), then in step 602, synchronization requests are sent from the initiating computer 10 to 1 - n target
5 computers 12. Next, each target computer 12 will respond to the request in step 604. Specifically, a target computer 12 can refuse the request such that synchronization does not occur. If the target computer 12 refuses, then in step 608, there is no further action with those target computers 12 which refuse synchronization. However, in step 610, target computers 12 which accept the synchronization request
10 proceed to receive changes as previously described for Figure 3.

[0035] If the user wishes to send and receive changes (i.e., indicating a bilateral exchange), then the initiating computer 10 will send the synchronization request to the 1 - n target computers 12 as shown in step 612. Next, each of the target
15 computers 12 will respond to the request by either refusing or accepting the request. If a target computer 12 refuses the request, then in step 618 no further action is taken with that target computer 12. However, if the target computer 12 accepts the request, then the bilateral synchronization process shown in Figure 5 commences and the target computer 12 will begin sending changes to the initiating computer 10
20 and then proceed to receive all changes therefrom, as shown in step 620.

[0036] Referring to Figure 6, if the user wishes just to receive changes only (i.e., unilateral exchange), then the initiating computer 10 will send the synchronization request to the target computers 12 in step 622. Next, the target computers 12 will
25 respond in step 624. If a target computer 12 refuses to send changes to the initiating computer 10, then no further action will occur, as shown in step 626. However, if the target computer 12 agrees to the synchronization process, then in step 628, the changes are sent to the initiating computer as described for Figure 2.

PATENT
Attorney Docket No. 033370-005

[0037] Additional modifications and improvements of the present invention may also be apparent to those of ordinary skill in the art. Thus, the particular combination of parts described and illustrated herein is intended to represent only a certain embodiment of the present invention, and not intended to serve as a
5 limitation of alternative devices within the spirit and scope of the invention.